



(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
13.11.2002 Bulletin 2002/46

(51) Int Cl.⁷: **G10L 15/18**

(21) Application number: 01401879.0

(22) Date of filing: 13.07.2001

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
 MC NL PT SE TR**
 Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Oudeyer, Pierre-Yves,
c/o Sony France S.A.
75004 Paris (FR)

(74) Representative: **Bertrand, Didier et al**
c/o S.A. FEDIT-LORIOT & AUTRES
CONSEILS EN PROPRIETE INDUSTRIELLE
38, Avenue Hoche
75008 Paris (FR)

(30) Priority: 11.05.2001 EP 01401203

(71) Applicant: **Sony France S.A.**
92110 Clichy (FR)

(54) Emotion recognition method and device

(57) The emotion recognition is performed by :

- extracting (70) a set comprising at least one feature derived from said signal, and
- processing (84, 90) the set of extracted feature(s) to detect an emotion therefrom.

There is further provided a step of low-pass filtering (74) the voice signal prior to extracting therefrom at least one feature of the set. The cut-off frequency for the low pass filtering is typically centred around 250 Hz.

The features are e.g. statistical quantities extracted from sampling a signal of the intensity or pitch of the voice signal.

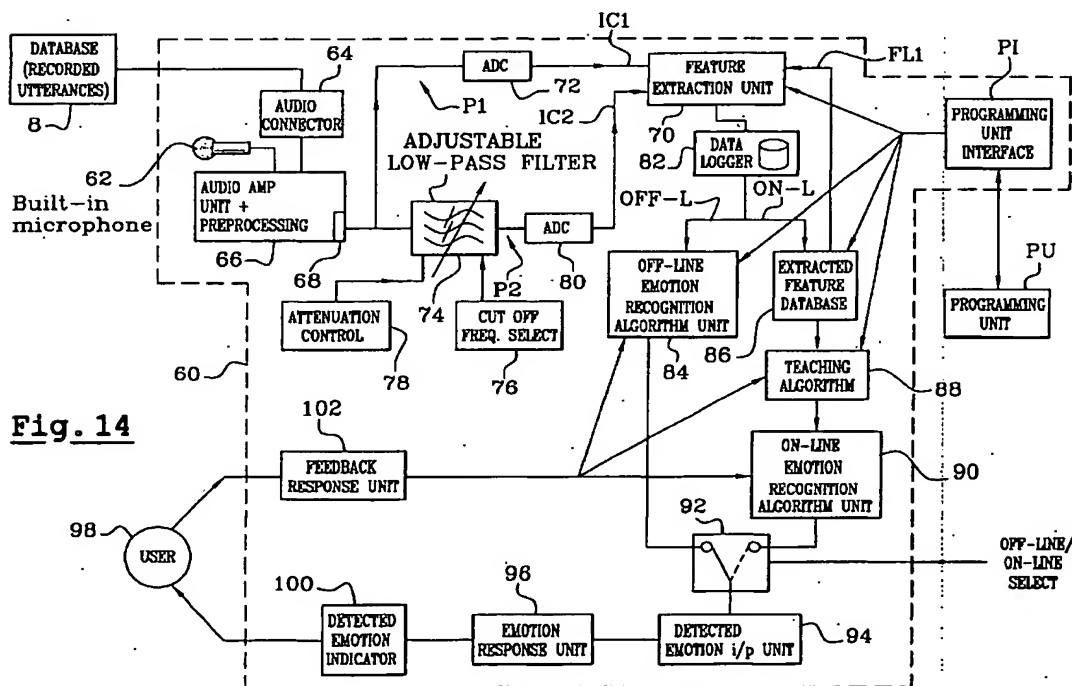


Fig. 14

Description

[0001] The present invention relates to the field of emotion recognition, and more particularly to a method and device for detecting an emotion from voice signals. There are now a growing number of applications which require technical means for detecting an emotion conveyed by a human - or possibly synthesised or animal - voice, e.g. to provide a man-machine interface with a more natural and intuitive way of communicating. For instance, robotic pets or humanoids are now developed with a behaviour remarkably similar to that of living creatures. In this context, it is natural that the robot be expected to detect utterances with some basic emotion content, e.g. calm, happiness, sadness, anger, etc. Other applications can be in education, training, speech therapy, human control of equipment, vehicles, tools, etc.

[0002] Most work (though not all) undertaken up to present in this area of emotion recognition concentrates on detecting emotions in a text read out by persons, the text generally being paragraphs rather than short utterances. The recitation is performed in well-controlled environments such as acoustic chambers.

[0003] Recent years have been marked by the increasing development of personal robots, either used as new educational technologies (cf. Druin A., Hendler J. (2000) "Robots for Kids: Exploring new technologies for learning", Morgan Kaufman Publishers.) or for pure entertainment (cf. Fujita M., Kitano H. (1998) "Development of an autonomous quadruped robot for robot entertainment", Autonomous Robots, 5, and Kusahara M. (2000) "The art of creating subjective reality: an analysis of Japanese digital pets", in Boudreau E., ed., in Artificial Life 7, Workshop Proceedings, pp. 141-144).

[0004] Typically, these robots look like familiar pets such as dogs or cats (e.g. the Sony AIBO robot), or sometimes take the form of young children such as the humanoids SDR3-X (Sony).

[0005] The interactions with these machines are to be radically different with the way we interact with traditional computers. So far, humans had the habit of learning to use very unnatural conventions and media such as keyboards or dialog windows, and had to have some significant knowledge about the way computers work to be able to use them.

[0006] Among the capabilities that these personal robots need, one of the most basic is the ability to have a grasp over human emotions (cf. Picard R. (1997) "Affective Computing", MIT Press), and in particular they should be able both to recognise human emotions and express their own emotions. Indeed, not only emotions are crucial to human reasoning, but they are central to social regulation (cf. Halliday M. (1975) "Learning how to mean: explorations in the development of language", Elsevier, NY.) and in particular to the control of dialog flows : orality and gestuality).

[0007] Apart from language, we express our emotions to others in two main ways: modulation of facial expression (cf. Ekman, P. (1982) "Emotions in the human face", Cambridge University Press, Cambridge.) and the modulation of the intonation of the voice (cf. Banse, R. and Sherer, K. R., (1996) "Acoustic profiles in vocal emotion expression", Journal of Personality and Social Psychology, 70(3): 614-636).

[0008] Whereas research on automated recognition of emotions in facial expressions is now very rich (cf. A. Samal, P. Iyengar (1992) "Automatic recognition and analysis of human faces and facial expression: A survey. Pattern Recognition, 25(1):65-77.), research on speech modality, both for automated production and recognition by machines, has been active only for very few years (cf. Bosh L.T. (2000) "Emotions: what is possible in the ASR framework ?", in Proceedings of the ISCA Workshop on Speech and Emotion.).

[0009] A starting point in the research is an investigation in acoustic correlates between emotion and its effect on the acoustic characteristics of the corresponding voice signal. A number of researchers have already investigated this question (cf. Fairbanks 1940, Burkhard and Sendmeier 2000, Banse and Sherer 1996).

[0010] Their results agree on the speech correlates that arise from physiological constraints, and which correspond to broad classes of basic emotions, but disagree and are unclear when one looks at the differences between the acoustic correlates of, for instance, fear and surprise, or boredom and sadness. Indeed, certain emotional states are often correlated with particular physiological states (cf. Picard 1997 cited supra) which, in turn, have quite mechanical - and thus predictable - effects on speech, especially on pitch (fundamental frequency F0), timing and voice quality. For instance, when one is in a state of anger, fear or joy, the sympathetic nervous system is aroused, heart rate and blood pressure increase, the mouth becomes dry and there are occasional muscle tremors. Speech is then loud, fast and enunciated with strong high frequency energy. When one is bored or sad, the parasympathetic nervous system is aroused, the heart rate and blood pressure decrease and salivation increases, producing speech that is slow, low-pitched and with little high frequency energy (Breazal, C. (2000) Sociable Machines : "Expressive social exchange between humans and robots", PhD Thesis, MIT AI Lab.).

[0011] Furthermore, the fact that these physiological effects are rather universal means that there are common tendencies in the acoustical correlates of basic emotions across different cultures. This has been precisely investigated in studies (cf. Abelin A., Allwood J., (2000) "Cross-linguistic interpretation of emotional prosody", in Proceedings of the ISCA Workshop on Speech and Emotion) or Tickle A. (2000), "English and Japanese speaker's emotion vocalisations and recognition : a comparison highlighting vowel quality", ISCA Workshop on Speech and Emotion, Belfast 2000). In these studies, experiments were conducted e.g. in which American people had to try to recognise the emotion of either another American or a Japanese person just by using acoustic information (the utterances were meaningless, so there

was no semantic information).

[0012] Japanese people were likewise asked to try to decide which emotions other Japanese or American people were trying to convey. Two results came out from these studies: 1) there was only little difference between the performance of trying to detect emotions conveyed by someone speaking the same language or the other language, and this held true for Japanese as well as for American subjects; 2) subjects were far from perfect recognisers in the absolute: the best recognition score was 60 percent. (This result could be partly explained by the fact that subjects were asked to utter nonsense utterances, which is quite unnatural, but is confirmed by studies asking people utter semantically neutral but meaningful sentences (cf. Burkhardt F., Sendlmeier W., (2000) "Verification of acoustical correlates of emotional speech using formant-synthesis", in Proceedings of the ISCA Workshop in Speech and Emotion.).

[0013] The first result indicates that the goal of making a machine express affect both with meaningless speech and in a way recognisable by people from different cultures with the accuracy of a human speaker is attainable in theory. The second result shows that we should not expect a perfect result, and compare the machine's performance in relation to human performance. The fact that humans are not so good is mainly explained by the fact that several emotional states have very similar physiological correlates and thus acoustic correlates. In actual situations, we solve the ambiguities by using the context and/or other modalities. Indeed, some experiments have shown that the multi-modal nature of the expression of affect can lead to a MacGurk effect for emotions (cf. Massaro D., (2000) "Multimodal emotion perception: analogous to speech processes", ISCA Workshop on Speech and Emotion, Belfast 2000.) and that different contexts may lead people to interpret the same intonation as expressing different emotions for each context (cf. Cauldwell R. (2000) "Where did the anger go? The role of context in interpreting emotions in speech", ISCA Workshop on Speech and Emotion.). These findings indicate that there is no need to have a machine generate utterances that make fine distinctions; only the most basic affects need be investigated.

[0014] A number of experiments using computer-based techniques of sound manipulation have been conducted to explore which particular aspects of speech reflect emotions with the most saliency (cf. Murray E. and Amott J.L. (1993) "Towards a simulation of emotion in synthetic speech", *Speech Communication*, 16(4), pp. 1097-1108; Banse R. and Sherer K.R. (1996) "Acoustic profiles in vocal emotion expression", *Journal of Personality and Social Psychology*, 70 (3): 614-636; Burkhardt and Sendlmeier (2000) "Verification of acoustical correlates of emotional speech using Formant-synthesis", Proceedings of the ISCA Workshop in Speech and Emotion; and Williams U. and Stevens K.N. (1972) "Emotions in speech: some acoustical correlates", *JASA* 52, 1228-1250. All these studies basically agree on the conclusion that the most crucial aspects are those related to prosody: pitch (on f_0) and contour, the intensity contour and the timing of utterances. Some more recent studies have shown that voice quality (cf. Gobl C. and Chasaide A.N. (2000) "Testing affective correlates of voice quality through analysis and resynthesis", Proceedings of the ISCA Workshop on Emotion and Speech) and certain co-articulatory phenomena (cf. Kienast M. and Sendlmeier W. (2000) "Acoustical analysis of spectral and temporal changes in emotional speech", Proceedings of the ISCA Workshop on Emotion and Speech) are also reasonably correlated with certain emotions.

[0015] Regarding of detection of human emotions, whereas humans generally do that by using all the context and modalities, ranging from linguistic content to facial expression and intonation, such a complete approach is unfortunately not realistic for a machine in an uncontrolled environment: for instance robust speech recognition in such situations is out of reach for present day systems, and facial expression recognition needs both computational resources and video devices that robotic creatures most often do not have.

[0016] For this reason, the Applicant investigated how emotion detection could be achieved by using only the prosodic information of the voice. Furthermore, the speech of interest is of the kind that occurs in everyday conversations, which means short informal utterances, as opposed to the speech produced when one is asked to read emotionally a prepared text, e.g. paragraph from a newspaper. Four broad classes of emotional content were studied: joy/pleasure, sorrow/sadness/grief, anger and calm/neutral.

[0017] As opposed to the automatic recognition of emotions with facial expression (cf. Samal A. and Iyengar P. "Automatic recognition and analysis of human faces and facial expression: a survey" *Pattern Recognition*, 25(1): 65-77), research using the speech modality is still very young (cf. Bosh L.T. (2000) "Emotions: what is possible in the ASR framework?" (Proceedings of the ISCA Workshop on Speech and Emotion, Belfast 2000). The first studies that were conducted (e.g. Murray and Amott (1993) "Towards a simulation of emotion in synthetic speech: a review on the literature of human vocal emotion", *JASA* 93(2), pp. 1097-1108, 1993; Williams U., Stevens K.N., (1972) "Emotions and speech: some acoustical correlates", *JASA* 52, 1238-1250) were not so much attempting to produce an efficient machine recognition device, but rather were searching for general qualitative acoustic correlates of emotion in speech (for example: happiness tends to make the mean pitch of utterances higher than in calm sentences). More recently, the increasing awareness of the fact that affective computing had an important industrial potential (Picard, 1997) pushed research towards the quest for performance in automatic recognition of emotions in speech (Bosh, 2000).

[0018] However, no large-scale study using the modern tools developed in the data mining and machine learning community have been conducted. Indeed, most often, either only one or two learning schemes are tested (for e.g. Polzin T., Waibel A. (2000) "Emotion-sensitive Human-computer Interface", in Proceedings of the ISCA Workshop on

Speech and Emotion; Slaney M., McRoberts G. (1998) "Baby Ears : a recognition system for affective vocalization, in Proceedings of ICASSP 1998), or very few and simple features are used (Polzin and Waibel 2000, Slaney and McRoberts 1998, Breazal 2000, Whiteside 1997), or only small databases are used - less than 100 examples per speaker (cf. C. Breazal; L. Aryananda (2000) "Recognition of affective communicative intent in robot-directed speech", in Proceedings of the Humanoids Conference, 2000; McGilloway S. et al. (2000) "Approaching automatic recognition of emotion from voice : a rough benchmark", in Proceedings of the ISCA Workshop on Speech and Emotion; Slaney M., McRoberts G. (1998) "Baby ears: a recognition system for affective vocalisation", in Proceedings of ICASSP 1998) which suggests that the power of some statistical learning schemes may have been overlooked.

[0019] Only McGilloway and al. 2000 cited supra have attempted to make some systematic data mining, using more than the traditional/standard set of features used by the rest of the literature : mean, max, min, max-min, variance of the pitch and intensity distributions, and of the lengths of phonemic or syllabic segments, or of pitch rising segments.

[0020] However, this work lacks an exhaustive experimental basis: 1) only three kinds of learning schemes were used - support vector machines, Gaussian mixtures and linear discriminants - which are far from being the best at dealing with data in which there are possibly many irrelevant features, and in particular which do not allow to derive automatically smaller set of features with optimal efficiency; and 2) the feature set was explored by choosing one learning scheme and iteratively removing less useful features for classification : on the one hand this is rather ad hoc since it is linked to a very particular learning scheme and selection procedure, and on the other it does not allow to detect the appropriateness of groups of features. Their work is based on speech generated by asking human subjects to read in an emotional manner newspaper texts, which does not correspond to constraints addressed by the present invention.

[0021] To this date, the only two research groups have apparently attempted to build automatic recognition machines of everyday speech are (Breazal 2000 and Slaney et al. 1998 cited supra). Yet, these groups could only use very small databases, very few simple features and two different learning algorithms. Finally, a general conclusion of this existing corpus of research is that recognition rates above 60 percent, even with only four basic emotions, seems impossible if there are several speakers to accommodate. The enormous speaker variability has for instance been described (cf. Slaney et al. 1998, cited supra).

[0022] It is therefore an object of the invention to propose a voice-based emotion recognition method and device which is affordable in terms of processing workload, so making it practical to implement in small systems such as pet robots, humanoids or the like, while offering a good success rate at correctly discerning an emotion.

[0023] Another object of the invention is to make it possible identify emotions conveyed in short utterances.

[0024] A further object of the invention is to allow operation in real life situations, where utterances are made in noisy and/or acoustically poor environments.

[0025] More particularly, the invention proposes, according to first aspect, a method of detecting an emotion conveyed in a voice signal, comprising the steps of:

- extracting a set comprising at least one feature derived from said signal, and
- processing said set of extracted feature(s) to detect an emotion therefrom,

characterised in that it comprises a step of low-pass filtering the voice signal prior to extracting therefrom at least one feature of that set.

[0026] The low-pass filtering step preferably comprises filtering the intensity of the voice signal with a cut-off frequency (Fco) substantially in the range of 150 to 400 Hz, more preferably substantially in the range of 250 to 300 Hz, a preferred cut-off frequency (Fco) being substantially equal to 250 Hz.

[0027] Advantageously, the feature extracting step comprises the steps of selectively extracting one or more among 160 features obtained from the following combination of parameters I, II, and III:

- Parameters I:

- i) the intensity of the signal derived after the low-pass filtering step,
- ii) the intensity of the signal without low-pass filtering,
- iii) the pitch of the signal without low-pass filtering, and
- iv) the pitch of the signal derived after the low-pass filtering step;

- Parameters II, from a sequence of values of a parameter I obtained over a duration of the voice signal:

- 1) a series of local minimas, whose members are values in the sequence that correspond to a minima,
- 2) a series of local maximas, whose members are values in the sequence that correspond to a maxima,
- 3) a series of durations (time intervals), whose members are time values separating local (adjacent) extremas,

and

4) a series, whose members are all the successive sampled values of the sequence; and

- Parameters III, calculated from a series of parameter II:

- 1) the mean,
- 2) the maximum (i.e. the overall highest value),
- 3) the minimum (i.e. the overall lowest value),
- 4) the difference between the maximum and the minimum,
- 5) the variance,
- 6) the median,
- 7) the first quartile,
- 8) the third quartile,
- 9) the interquartile range, defined as the difference between the third and first quartiles,
- 10) the mean of the absolute value of the local derivative;

at least one of the features originating from: i) the intensity of the signal after the low-pass filtering step or from iv) the pitch of the signal after the low-pass filtering step of parameters I.

[0028] Preferably, the parameters I comprise more particularly:

- i) the intensity of the signal derived after the low-pass filtering step,
- ii) the intensity of the signal without low-pass filtering, and
- iii) the pitch of the signal without low-pass filtering,

whereby the combination yields a total of 120 different possible features;

at least one of the features originating from: i) the intensity of the signal after the low-pass filtering step of parameters I.

[0029] Preferably, the parameters III comprise more particularly the following five parameters:

- 2) the maximum (i.e. the overall highest value),
- 3) the minimum (i.e. the overall lowest value),
- 6) the median,
- 7) the first quartile,
- 8) the third quartile,

[0030] Advantageously the extracting step comprises extracting at least one of the following possible features:

- the maximum (i.e. the overall highest value),
 - the minimum (i.e. the overall lowest value),
 - the median,
 - the first quartile,
 - the third quartile,
- calculated from:

- the series of values of the intensity signal derived after the low-pass filtering of the voice signal,
- the series of minimum values of the intensity signal derived without low-pass filtering the voice signal,
- the series of pitch values of the signal derived after the low-pass filtering of the voice signal, and
- the series of pitch values of the signal derived without low-pass filtering of the voice signal.

[0031] The extracting step can comprise extracting features derived from both the voice signal after the low low-pass filtering step and from the voice signal without low-pass filtering.

[0032] It may also comprises extracting features only from the voice signal after the low-pass filtering step.

[0033] Advantageously, the at least one member of the set of features comprises:

- 1: the median value of the signal intensity values derived after low-pass filtering the voice signal,
- 2: the mean value of the signal intensity values derived after low-pass filtering the voice signal,
- 3: third quartile value of the signal intensity values derived after low-pass filtering the voice signal,
- 4: the first quartile value of the signal intensity values derived after low-pass filtering the voice signal,

EP 1 256 937 A2

- 5: the maximum of the signal intensity values derived after low-pass filtering the voice signal, and
6: the maximum value of the series of the signal intensity values derived after low-pass filtering the voice signal.

5 [0034] The method may further comprise a step of smoothing a signal extracted from the voice signal prior to the extracting step.

[0035] In an embodiment, the method further comprises a step of generating an emotion detecting algorithm using a teaching algorithm, the teaching algorithm exploiting the at least one feature extracted from a low-pass filtered voice signal.

10 [0036] The method may further comprise a step of programming an emotion detection device with an off-line emotion detection algorithm and/or an on-line emotion detection algorithm, whereby the algorithms exploit at least one feature extracted from a low-pass filtered voice signal.

[0037] According to a second aspect, the invention concerns a device for detecting an emotion conveyed in a voice signal, comprising the:

- 15 - extracting means for extracting a set comprising at least one feature derived from the signal, and
- processing means for processing the set of extracted feature(s) to detect an emotion therefrom,

characterised in that it comprises a low-pass filtering means for low-pass filtering the voice signal prior to extracting therefrom at least one the feature of the set.

20 [0038] The optional features presented above in the context of the method apply mutatis mutandis to the above device, and shall not be repeated for conciseness.

[0039] According to a third aspect, the invention relates to a voice signal processing device specifically for the device according to the second aspect - or specifically adapted for the implementation of the method according to the first aspect -, characterised in that comprises low-pass filtering means for supplying a low-passed form of the voice signal to the extracting means.

25 [0040] The invention and its advantages shall become more apparent from reading the following description of the preferred embodiments, given purely by way non-limiting examples with reference to the appended drawings in which :

- 30 - figure 1 is a chart showing how some basic emotions are linked to such parameters as valence and excitement;
- figures 2a, 2b and 2c are signal waveforms showing the relation between a speech signal (figure 2a), pitch (figure 2b) and intensity (figure 2c);
- figure 3 is a block diagram of a system for evaluating, characterising and optimising an emotion detection approach using low-pass filtering of the audio signal on which an emotion is to be detected, in accordance with the invention;
- figure 4 is a curve showing a typical and an ideal low pass filter response curve;
35 - figure 5 is a set of curves showing the response of an adjustable cut-off frequency filter used in the system of figure 3;
- figure 6 is a sample signal waveform (signal intensity vs. time) showing how different signal features considered are derived;
- figures 7a, 7b and 7c are plots showing different distributions of voice samples when analysed in terms of different extracted signal features for emotion detection, respectively for a poorly informative, good and even better chosen feature;
40 - figures 8a and 8b are plots of detected emotion in terms of first and third quartile intensity distribution expressing the degree of success in emotion detection in terms of grouping characteristics of data points, respectively for a detection on an unfiltered signal and for a low-pass filtered signal with a cut-off frequency of 250 Hz;
- figures 9 to 13 are charts showing the performance of different algorithms used for emotion detection under different conditions and choices of signal features; and
45 - figure 14 is a block diagram of an emotion sensitive system implementing a low-pass filtering of voice signals in accordance with the invention.

50 [0041] The preferred embodiment of the invention is chosen to be speaker dependant for the emotion recognition process. This is not necessarily a bad point from an industrial point of view, when the embodiment it is targeted at robotic pets, for instance, which may interact only with their minder. Also, the fact that robots only manage to recognise their minder could even be a positive feature, because it is a source of complicity between a robot and its minder). However, the principles of invention can of course accommodate for speaker-independent approaches, as shall be outlined further.

55 [0042] The embodiment is based on a robotic pet application, with four basic emotional states to detect in a human voice : i) calm, ii) sad, iii) angry and iv) happy.

[0043] Figure 1 shows how these emotions are positioned in a chart which represents an "emotional space", in which the parameters "valence" and "excitement" are expressed respectively along vertical and horizontal axes 2 and 4. The

valence axis ranges from negative to positive values, while the excitement axis ranges from low to high values. The cross-point O of these axes is at the centre of the chart and corresponds to a calm/neutral state. From that point are defined four quadrants, each containing an emotional state, as follows : happy/praising (quadrant Q1) characterised by positive valence and high excitement, comfort/soothing (quadrant Q2) characterised by positive valence and low excitement, sad (quadrant Q3) characterised by negative valence and low excitement, and angry/admonishing (quadrant Q4) characterised by negative valence and high excitement.

[0044] The constraints imposed on the embodiment were as follows :

- persons may talk very freely to the robotic pet,
- persons may utter very short sentences, such as "Aibo ! you're a really naughty robot!",
- the utterances will have to be detected in a noisy environment, and
- several persons may have to interact with the robot, whereupon a choice of speaker dependency or independence needs to be made.

[0045] Some of these constraints implied that the utterances be restricted to infant-directed type of speech.

[0046] At an initial stage, a strategy for study was developed to determine how best to meet these objectives. This strategy can be summarised as follows:

- i) experiment with simple features, simple algorithms and small databases,
- ii) explore systematically features and learning algorithms for the speaker-dependent case:

- a) with many examples, using a database with samples from two speakers, one male, the other female,
- b) with only a few examples of the above two-speaker database, and

- iii) conduct experiments in real-life situations with fluent speech.

Database for study

[0047] The study was conducted using very large database of two speakers containing informal short emotional utterances. All experiments were conducted using the freely available data mining software Weka which implements most of the standards data mining techniques, information on which can be found in the Weka web page at: <http://www.cs.waikato.ac.nz/~ml/>.

[0048] In order to have sufficiently large databases, some compromises had to be made (the recording conditions as described in Slaney et al, 1998 or Breazal 2000 are rather poor and impractical). Accordingly, the database used for the experimental study used two Japanese professional speakers (respectively male and female), who are both voice actor/actress and worked on many radio/TV commercials, Japanese dubbing of movies and animations. They were asked to imitate everyday speech by pronouncing short sentences or phrases like "Umm, I don't know", "Exactly!", "See", "Hello", "I see", "How are you?", "What kind of food do you like?", "Wonderful!", "D'you know". Before each utterance, they had to imagine themselves in a situation where they could produce the phrase, and which would correspond to one of the four emotional classes: joy/pleasure, sorrow/sadness/grief, anger, normal/neutral. If several emotions were compatible with the sentence meaning, then they were allowed to utter each phrase with the corresponding emotion. The resulting database comprised 200 examples per speaker and per emotion, making a total of 2000 samples.

[0049] The database stored in digital form the sound recordings of the samples, so that each sample could be accessed and played back at will.

[0050] The two main measurements that were undertaken concerning intonation were pitch and intensity, as in all the reported works reported mentioned supra.

[0051] Figures 2a-2c illustrate how the two physical parameters pitch and intensity are derived from a speech signal. The latter is illustrated in figure 2a as a waveform envelope, as obtained e.g. directly from a microphone output or from playing back a voice sample. The waveform is the fluctuation in signal amplitude (vertical axis) as a function of time (horizontal axis). In the representation, the signal strength is assigned a zero value at a median position, from which it can take on negative (low amplitude) and positive (high amplitude) values.

[0052] Pitch is a subjective quality of a sound that determines its position in a musical scale or voice spectrum. It may be measured as the frequency of the pure tone of specified intensity that is judged by the average ear to occupy the same place in the musical scale or voice spectrum. Figure 2b illustrates a variation of pitch with time. Here, pitch is measured in terms of a frequency (vertical axis) for a particular time value.

[0053] Intensity is a measure of the sound energy integrated over a given frequency spectrum. Figure 2c illustrates the variation of intensity with time, with the intensity measured in terms of decibels (dB).

[0054] Additionally, in accordance with a feature specific to the present invention, measurements were made on a low-pass filtered form of the electronic signal conveying the utterances, hereafter referred to as the voice signal.

[0055] Figure 3 illustrates schematically a system 1 for evaluating emotion detection performance with the possibility of low-pass filtering the voice signal.

5 [0056] The voice signal can be obtained either directly from a microphone 6 or from the above-mentioned database of recorded utterances 8. In the latter case, each of the samples can be selectively read out in audio format, e.g. from a hard disk drive, so as to obtain an audio signal of the corresponding utterance at an audio output 10. The signal from the microphone 6 or database audio output 10 is fed to an audio input 12 of an audio amplification unit 14 which provides impedance matching, preamplification, and controllable gain functions. The thus-prepared audio signal is delivered at an audio output 14, from which it is fed to a low-pass filter 18. In the system 1, the low-pass filter 18 has an adjustable cut-off frequency F_{co} , the value of which is selected by a cut-off frequency selection unit 20.

10 [0057] Figure 4 shows a typical frequency response of the low-pass filter 18 with a curve 22 showing the variation of transferred audio signal level against frequency for an input of constant amplitude over the frequency range. The cut-off frequency F_{co} is shown as the point of half attenuation along the frequency scale. Ideally, the cut-off frequency is the frequency value beyond which the low-pass filter does not transfer a signal. The sharpness depends on the filter characteristics, e.g. the order of filtering, to approach the ideal low-pass response 24.

15 [0058] Figure 5 shows how the frequency response changes according to the cut-off frequency F_{co} selected through the selection unit 20, with three curves illustrated 22a, 22b and 22c corresponding to F_{co} values respectively of 200 Hz, 250 Hz and 320 Hz. The shape of the response curves are substantially the same. In the system of figure 3, the low-pass filter has a continually adjustable cut-off frequency F_{co} in view of obtaining fine level comparisons. The cut-off frequency selection unit 20 can be electronically controlled to vary a filter component, such as one or several varicap diodes.

20 [0059] The filter 18 is constructed as an analog circuit based on operational amplifiers. Such adjustable cut-off frequency low-pass filters are well known in the art and shall not be described here for conciseness. It can of course equally well be envisaged to produce the low-pass filtering function entirely by digital filtering techniques using e.g. digital signal processors, which are also well known in the art. In this case, the input audio signal would be initially digitised for processing by the filter electronics, the latter reproducing the response curve 22 in a suitably digitally encoded form.

25 [0060] The output from the low-pass filter 18 is then supplied to a first input 28a of an audio interface 28 of a feature measurement station 30, which is represented as a functional unit cooperating with a PC based control unit 32.

30 [0061] There is also a separate leading from the audio output 16 directly to a second input 28b of the audio interface 28 of the feature measurement station 30. The two inputs 28a and 28b constitute separate input channels allowing the feature measurement station 30 to operate selectively on the low-pass filtered form and non filtered forms of the audio signal.

35 [0062] The PC based control unit 32 is composed of a monitor 34 and keyboard 36 forming the system interface, and a complete PC type computer 38 for the processing, drive and data storage functions. The feature measurement station 30 is programmable to extract different features of the signal presented at its signal input 28. As feature extraction provided by the station 30 is based on digital processing techniques, its signal input 28 is connected to an internal analog-to-digital converter (ADC) 40. The filtered audio signal is then processed in terms of binary words to extract the required features, which are discussed below.

40 [0063] The measurements on the different features produced by the station 30 are stored in a data logger 42 where they can be accessed by a functional unit 44 which performs emotion recognition algorithms. The algorithms aim to discern an emotion expressed in the voice signal on the basis of the signal features extracted by the station 30. Typically, an algorithm under evaluation in unit 44 will indicate the discerned emotion (calm, happy, ...) and various internal parameters. The recognition unit 44 can thus execute different emotion recognition algorithms and produce their respective results and internal parameters.

45 [0064] This information is entered into a results analyser unit 46 where different algorithms and/or choices of features can be compared in view of determining an optimum selection.

50 [0065] It will be appreciated the functional units composed of the feature measurement station 30, data logger 42, emotion recognition algorithm execution unit 44 and the results analyser 46 can be implemented physically by the PC based control unit 32, as indicated by the staggered double-arrows.

55 [0066] The system further includes a programming unit PU whose function is to program different emotion recognition systems on the basis of the most appropriate results found. To this end, the programming unit PU is operationally connected to the results analyser 46 as well as to the PC based control unit 32 and other units, so as to record all the algorithms, parameter settings, features to extract and other optimisations. The recording of this information is established in accordance with a predetermined protocol so that the programming data can be downloaded into a suitable programming interface of the emotion sensitive system, such as a robotic pet or the like. The programming unit PU can be made detachable from the rest of the system and portable for on-field programming, and/or adapted to produce

the programming data on a portable recording medium, such as a disk, semiconductor memory, or the like.

[0067] The measurements performed by the feature measurement station 30 were based on the intensity of low-passed form of the voice signal, as obtained from the filter 18. (Separate experiments were also conducted on high-passed forms using a high-pass filter instead.)

[0068] For the low-pass filtering, the cut-off frequency was typically set to 250 Hz (the particular value appears not to be crucial), the idea being to separate the signal into a pure prosodic component and a pure "cepstral" component.

[0069] The samples from the database 8 thus provide sound patterns on which various measurements can be made at the station 30. Psycho-acoustic measurements have shown that the most important information is prosody (intonation), and in particular the pitch curve (f0) (cf. figure 2b) and the intensity (I0) curve (cf. figure 2c).

[0070] In the evaluation, three different signals were processed by the station 30 :

- i) the intensity of the signal after low-pass filtering,
- ii) the intensity of the signal without low-pass filtering, and
- iii) the pitch of the signal without low-pass filtering.

[0071] The first signal i) is taken from input 28a, while the last two signal is taken from input 28b of the feature measurement station 30.

[0072] Note that the pitch can also be measured on the signal after low-pass filtering either instead of on the non filtered signal iii), or in addition, in which case the number of signals in the evaluation total four. However, experiments have shown that the information provided by the pitch analysed from the low-pass filtered signal is generally not significantly different from that provided by the pitch of the unfiltered signal. Accordingly, only signals i) to iii) are actually used in the evaluation.

[0073] After low pass filtering, the intensity curve of the resultant signal i) was measured in the same way as for the unfiltered signal.

[0074] Finally, to be exhaustive, there was conducted a cepstral measurement on signals iv) computed from the norm of the absolute vector derivative of the first 10 MFCCs (mel-frequency cepstral components). In speech recognition, there is often performed a measurement at each time step of the cepstral content, which is a parameter equivalent to the excitation of different nerves in the cochlea. These can correspond to the different cepstral coefficients of the signal. The measurement establishes the mean of the local slope of the derivatives of the vectors in the cepters. Typically, these measurements yield ten coefficients with ten values that evolve in time, and establish the difference between the successive vectors, from which can be obtained the norm of the vectors.

[0075] All these measurements i) to iv) above were performed at each 0.01s time frame, using a signal processing toolkit freely available under the name "Praat software", information on which can be found in web page " <http://www.praat.org>".

[0076] In particular, the pitch was computed using the algorithm described in Boersma P. (1993) "Accurate short-term analysis of the fundamental frequency and the harmonics-to-noise ratio of a sampled sound", in Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam, 17, 97-110 1993, which is known to be very accurate.

[0077] In what follows, the cepstral measurement signals (iv) are taken into account in the evaluation; only measurements on signals i) to iii) cited above shall be considered.

[0078] Each of these measurements provides a series of values that were transformed to produce different points of view on the data.

[0079] The procedure to obtain the series of values is as follows. For a given voice sample from the database 8, typically corresponding to a short phrase, the feature measurement station 30 stores the corresponding voice signal produced at the input (channel 28a or 28b depending on the signal i), ii) or iii) above to be analysed). The signals are initially 10 Hz smoothed.

[0080] From that stored 10 Hz smoothed signal are computed each of the three time-varying signals i) to iii) mentioned above.

[0081] Each thus-computed signal i) to iii) is digitally sampled at e.g. 0.01 second intervals. Each sampling gives an absolute instantaneous value of the intensity (signals i) and ii)) or of the pitch (signal iii)).

[0082] For each sampled signal, there is thus obtained a sequence of sampled values, typically a few hundred, given that a phrase is a few seconds long. The sequence of values is stored in time order, starting from the value of the earliest part of the signal, so that the sequence of values expresses the time variation of signal i), ii) or iii) considered.

[0083] For each signal, the corresponding time-ordered sequence of values is analysed to produce the following four series of values :

- 1) The series of local minimas. This involves determining the values in the sequence that correspond to a minima and constructing a series formed just of these minima values of the sequence. The series is stored in time order, starting from the earliest minima. A minima in the sequence of values is identified as a member of the sequence

whose value is less than that of its neighbour immediately before and immediately after in the sequence.

2) The series of local maximas. This involves determining the values in the sequence that correspond to a maxima and constructing a series formed just of these maxima values of the sequence. The series is stored in time order, starting from the earliest maxima. A maxima in the sequence of values is identified as a member of the sequence whose value is greater than that of its neighbour immediately before and immediately after in the sequence.

3) The series of durations (time intervals) between local (adjacent) extremas (which models rhythmic aspects of the signal). An extrema is either a minima or a maxima. This involves first identifying in the sequence the values that correspond to a minima or a maxima, by comparing each value with its immediately adjacent neighbours as for series 1) and 2) above. Then, the positions in time of the identified extremas are analysed to determine the time intervals separating successive extrema. These time intervals are set out in a series, starting from the interval separating the earliest and next earliest extremas, to yield the series of time intervals.

4) The series itself. This simply involves forming a series from all the successive sampled values of the sequence, starting from the beginning of the sampled voice signal.

[0084] Figure 6 illustrates how these four series of values are obtained from the signal undergoing measurement.

[0085] It will be noted that for each signal i) to iii) above measured, the four series of values 1) to 4) above are extracted, giving a total of 12 series of values per audio sample of the database.

[0086] Finally, the following ten statistical characteristics are computed for each of the twelve series of values:

- 1) the mean,
- 2) the maximum (i.e. the overall highest value),
- 3) the minimum (i.e. the overall lowest value),
- 4) the difference between the maximum and the minimum,
- 5) the variance,
- 6) the median,
- 7) the first quartile,
- 8) the third quartile,
- 9) the interquartile range, defined as the difference between the third and first quartiles,
- 10) the mean of the absolute value of the local derivative.

[0087] All the data are initially normalised before calculating these statistical characteristics.

[0088] The resulting combination of these ten characteristics applied to the four series gives rise to $3 \times 4 \times 10 = 120$ features.

Learning algorithms

[0089] Next, after all these values are obtained for each signal, an experiment was conducted using different automated learning methods. The latter are capable of learning and generalising from a certain number of examples. Many such methods are known, each having their advantages and disadvantages. The approach makes a fairly exhaustive use of these different algorithms.

[0090] For each learning algorithm tested, there is input a base set of examples. The performance for each base set is then checked against another example base set. This allows to determine a percentage of correct generalisations made by each learning algorithm.

[0091] There are many learning schemes that have been developed in the last 20 years (e.g. Witten I., Frank E. (2000) "Data Mining", Morgan Kaufman Publishers.), and they are often not equivalent : some are more efficient with certain types of class distributions than others, while some are better at dealing with many irrelevant features (which is the case here, as seen a posteriori) or with structured feature sets (in which this is the "syntactic" combination of the values of features which is crucial). As, by definition, the structure of the data and/or the (ir-)relevance of its features is unknown, it would be a mistake to investigate the problem with only very few learning schemes. As a consequence, it was chosen to use a set of the most representative learning schemes, ranging from neural networks to rule induction or classification by regression. Also, one of the best meta-learning scheme, was used, namely AdaBoostM1 (Witten and Frank, 2000 cited supra), which significantly improves the generalisation performance for unstable learning schemes like decision trees (an unstable learning algorithm is one that can sometimes produce very different recognition machines when only a slight change in the learning database has been performed). One reason for choosing the Weka software was that its code and executable are freely available, so that the experiment, though being large scale, can be easily reproduced. This software also provides means like automatic cross-validation, or the search of feature spaces with for e.g. genetic algorithms as shall appear later. The list of all learning algorithms is given in table I below.

Table I

list of learning algorithms investigated	
Name	Description
1-NN	1 nearest neighbour
3-NN	voted two nearest neighbours
10-NN	voted ten nearest neighbours
Decision tree-C4.5	C4.5 decision trees
Decision Rules-PART	PART decision rules
Kernel Density	Radial basis function - neural network
KStan	--
Linear Regression	Classification via linear regression
LWR	Classification via locally weighted regression
Voted Perceptrons	Committee of perceptrons
SVM1	Polynomial (deg.1) support vector machine
SVM2	Polynomial (deg.2) support vector machine
SVM3	Polynomial (deg.3) support vector machine
VF1	Voted features interval
M5Prime	Classification via MP5Prim regression method
Naive Bayes	Naive Bayes classification algorithm
AidaBoots M1/C4.5	Ada boosted version of C4.5
AidaBoots M1/PART	Ada boosted version of PART

[0092] More details about these algorithms can be found in Witten and Frank, 2000 cited supra.

[0093] In a first experiment, an evaluation was conducted in which all algorithms were attributed all the (normalised) features, and were trained on 90 percent of the database and tested on the remaining 10 percent. This was repeated 10 times, with each time a different 90/10 percent split (a 10-fold cross-validation thus being performed).

[0094] Table II below gives the average percentage of correct classification for the 10 fold cross-validation for the different algorithms used.

Table II

results using all features for two different speakers (1 and 2)		
Algorithm	Emotion detection success rate (%)	
	Speaker 1	Speaker 2
1-NN	82	87
3-NN	84	87
10-NN	83	87
Decision tree-C4.5	84	93
Decision Rules-PART	84	94
Kernel Density	84	90
KStan	81	83
Linear Regression	88	91
LWR	87	90

Table II (continued)

results using all features for two different speakers (1 and 2)		
Algorithm	Emotion detection success rate (%)	
	Speaker 1	Speaker 2
Voted Perceptrons	70	76
SVM1	88	94
SVM2	89	94
SVM3	88	94
VFI	80	93
M5Prime	86	96
Naive Bayes	84	90
AidaBoots M1/C4.590	90	96
AidaBoots M1/PART	91	97

[0095] It can be observed from these results that very high success rate (between 92 and 97 percent), which is higher than any other reported result in the literature can be obtained thanks to the use of certain algorithms. However, the difference among algorithms is striking: whereas the best results are obtained with "Adaboosted" decision trees and rules, some others perform 10 percent below (like nearest neighbours, RBF neural nets or Support Vector Machines, which are the ones typically used in other studies), or even 20 percent below (Committees of perceptrons). This illustrates that care must be taken to try many different learning schemes when solving a problem about which there is very little by way of few prior or intuitive knowledge.

Feature selection

[0096] After this first experiment, the next step is to see how the feature set could be reduced, for three reasons :

- 1) a small features set can be expected to provide better generalisation performance (see Witten and Frank, 2000 cited supra);
- 2) it is computationally cheaper to compute fewer features;
- 3) it is interesting to see if the most useful features for the machine learning algorithms are the ones that are traditionally put forward in the psychoacoustic literature.

[0097] A first way of exploring the feature set is to look at the results of learning schemes like decision rules (PART), which are often used mainly as knowledge discovery devices.

[0098] One specific choice of rules and parameterisation produced and tested by the Applicant is:

Example 1**[0099]**

5 If MEDIANINTENSITYLOW > 0.48 and
 MINMINIMASPITCH <= 0.07 and
 10 THIRDQUARTINTENSITY > 0.42 ==> CALM
 ELSE If MEANINTENSITYLOW <= 0.58 and
 MEDIANINTENSITYLOW <= 0.29 ==> ANGRY
 15 ELSE If THIRDQUARTINTENSITYLOW > 0.48 ==> SAD
 ELSE ==> HAPPY

(the symbol "<=" means "less than or equal to").

20 **[0100]** All the numerical values are normalised.

[0101] The definition of declared variables in the above algorithm are as follows :

- MEDIANINTENSITYLOW : the median value of series of values of the intensity of the signal after low-pass filtering;
- MINMINIMASPITCH : the value smallest minima in the series of minima in the pitch of the signal without low-pass filtering;
- 25 - THIRDQUARTINTENSITY : the value of the third quartile in the series of values of the intensity of the signal without low-pass filtering;
- MEANINTENSITYLOW : the mean value of the of the series of intensity values of the signal after low-pass filtering;
- THIRDQUARTINTENSITYLOW : the value of the third quartile in the series of intensity values of the signal after low-pass filtering.

[0102] These four and surprisingly simple rules allow a remarkably high percentage of correct classification in generalisation of 94 percent for the speaker 2 database. The striking fact is the repeated use of features related to the intensity of the low-pass signal.

35 **[0103]** Different sets rules can be adapted for optimised performance with different speakers/languages. The choice of an appropriate set of rules for a given speaker/language can be determined empirically.

[0104] Another view of the feature set can be obtained simply by visualisation. An example of such a visualisation is illustrated in figures 7a-7c, which each depict an example of a line of measurements for a feature i, j and k respectively taken among the 160 features investigated, as explained above. For each line is defined a space in a rectangular box 48a-48c within which are plotted a number of points. Each plotted point corresponds to an identification of the voice samples of the database 8. The position of a point corresponds to the associated measurement value for the feature considered (evolving on the horizontal scale). The emotion corresponding to the plotted samples is identified by the following key : dot = calm, cross = sad, dash = angry, circle = happy. In the complete experiment, 160 such lines are produced, one for each feature.

45 **[0105]** Figure 7a shows a typical haphazard distribution of identified emotion samples, showing that the feature i from which it is generated does not provide a good basis for discerning emotions. Such a distribution is frequent among features extracted from a voice signal that has not been low-pass filtered.

[0106] Figure 7b shows a significantly more clustered distribution of identified emotion samples, indicating that the feature from which it is generated is potentially good for discerning emotions. However, there is some overlap between neighbouring clusters. Such a distribution is quite typical of a feature extracted from a voice signal whose intensity has been low-pass filtered, the cut-off frequency (f0) being 250 Hz in the examples.

[0107] Finally, figure 7c shows an even better distribution of identified emotion samples where the emotion samples are not only well clustered, but also clearly separated by empty regions. Such results are obtained with some features extracted from a voice signal whose intensity has been low-pass filtered.

55 **[0108]** (Note : to avoid cluttering, only a representative fraction of the voice samples are plotted in the figure.)

[0109] The impressive results obtained by low-pass filtering the intensity of the voice signal was confirmed with different speakers.

[0110] In the experiment, six features were selected among the group, for having the best clustering of emotion

samples with at least some empty regions separating them.

[0111] Generally, the best results are obtained by considering the following features of the low-passed filtered intensity of the voice signals :

- quartiles, notably the first and third quartiles,
- maximas,
- minimas, and
- the mean.

[0112] By contrast, the least conclusive results were obtained from features involving the variance or mean of the absolute slope of the low-pass filtered signal.

[0113] To confirm the Applicant's intuition that low-passed intensity is crucial in the distinction of emotions, figures 8a and 8b show plots of identified emotions from the samples of the database 8. In these plots, two features noted for yielding potentially good emotion discernment are expressed on respective orthogonal axes: the first quartile of intensity distribution on the vertical (ordinate) axis, and the third quartile of intensity distribution on the horizontal axis (abscissa). The symbols used to identify the four different emotions are the same as in figures 7a-7c. Similarly, only a representative fraction of symbols are inserted to avoid cluttering.

[0114] The plots of figure 8a are produced from the voice samples of the database without low-pass filtering, while the plots of figure 8b are produced from the same voice samples, but after low-pass filtering with a cut-off frequency of 250 Hz using the filter 18 of figure 3. The voice samples in both cases came from the same speaker (speaker 2). The same very striking effect also takes place for speaker 1; but what is interesting is that the clusters are not situated at the same places (anger and happiness are 90 degrees rotated), which is an illustration of the great speaker variability initially suspected. However, it seems that the use of some well chosen features is very stable for each speaker.

[0115] In order to quantify the individual relevance of features or attributes, there is a measure that is often used in the data mining literature, which is the expected information gain, or mutual information between class and attribute. It corresponds to the difference between the entropies $H(\text{class})$ and $H(\text{class}|\text{attribute})$ (see Witten and Frank, 2000 cited supra, for details about how it is computed). Table III below gives the 20 best attributes according to the information gain they provide.

Table III:

Information gain of 20 best features	
Feature	Information gain
1: MEDIANINTENSITYLOW	1.44
(median value of sampled low-pass filtered signal intensity values)	
2: MEANINTENSITYLOW	1.40
(mean value of low-pass filtered signal intensity values)	
3: THIRDQUARTINTENSITYLOW	1.35
(third quartile value of series of sampled low-pass filtered signal intensity values)	
4: ONEQUARTINTENSITYLOW	1.34
(first quartile value of series of sampled low-pass filtered signal intensity values)	
5: MAXINTENSITYLOW	1.23
(maximum value of series of sampled low-pass filtered signal intensity values)	
6: MININTENSITYLOW	1.14
(maximum value of series of sampled low-pass filtered signal intensity values)	
7: THIRDQUARTMINIMASPITCH	0.72
(maximum value of series of minimas of low-pass filtered signal intensity values)	
8: THIRDQUARTMAXIMASPITCH	0.72
(maximum value of series of maximas of low-pass filtered signal intensity values)	
9: THIRDQUARTPITCH	0.69

Table III: (continued)

Information gain of 20 best features	
Feature	Information gain
(third quartile of series of sampled unfiltered pitch values)	
10: MAXMINIMASPITCH	0.67
(largest value in series of minimas of sampled unfiltered pitch values)	
11: MAXMINI	0.67
(largest value in series of maximas of sampled unfiltered pitch values)	
12: MAXPITCH	0.67
(maximum value of sampled unfiltered pitch values)	
13: MINMINIMASPITCH	0.59
(smallest value in series of minimas of sampled unfiltered pitch values)	
14: MEDIANMAXIMASPITCH	0.57
(median value in series of minimas of sampled unfiltered pitch values)	
15: MEDIANMAXIMASPITCH	0.57
(smallest value in series of minimas of sampled unfiltered pitch values)	
16: MINPITCH	0.52
(minimum value of series of sampled unfiltered pitch values)	
17: MEDIANPITCH	0.52
(median value of series of sampled unfiltered pitch values)	
18: MEANMINIMASPITCH	0.48
(mean value in series of minimas of sampled unfiltered pitch values)	
19: MINMAXIMASPITCH	0.48
(mean value in series of maximas of sampled unfiltered pitch values)	
20: MEANPITCH	0.48
(mean value of series of sampled unfiltered pitch values)	

[0116] This table confirms the great value of the features concerning the first six set of the distribution of intensity values involving the low-passed signals. It also shows something rather surprising: among the 20 most individually informative features, only 3 (the features numbered 12, 16 and 20 in the table) are part of the standard set put forward in psychoacoustic studies (Murray E., Arnott J.L., (1995) "Implementation and testing of a system for producing emotion-by-rule in synthetic speech", Speech Communication, 16(4), pp. 369-390; Sendmeier and Burkhardt 2000 cited supra; Stevens and Williams 1972 cited supra), or used in most of more application oriented research as in Slaney et al. 1998 cited supra, Breazal 2000 cited supra).

[0117] Nevertheless, it must be noted that individual salience of a feature is only partially interesting: it is not rare that success comes from the combination of features. Accordingly, in a first experiment, the Applicant compared a feature set containing only the features 1 to 6 related to low-passed signal intensity (LPF), with a feature set composed of the standard features (SF) used in (Breazal 2000, or Slaney et al. 1998) :

- 1) mean,
- 2) min,
- 3) max,
- 4) max-min,
- 5) variance of pitch of the unfiltered signal,
- 6) intensity of unfiltered signal, and
- 7) mean length of syllabic segments.

(Results are similar if jitter and tremor are added, as sometimes also used).

[0118] Table IV summarises these experiments (each number corresponds again to the mean percentage of correct classification in generalisation in 10-fold cross-validation).

Table IV

comparison between "standard" features and "low-passed signal intensity" features				
Learning schemes	(LPF)sp.1	(LPF)sp.2	(LPF)sp.1	(LPF)sp.2
1-NN	78	83	70	72
3-NN	84	82	72	75
10-NN	84	82	73	73
Decision Tree-C4.5	80	84	72	81
Decision Rules-PART	78	83	72	74
Kernel Density	82	85	71	74
Kstar	80	84	70	72
Linear Regression	63	68	72	78
LWR	75	71	75	80
Voted Perceptrons	51	70	60	58
SVM degree 1	63	68	73	78
SVM degree 2	71	70	77	50
SVM degree 3	76	85	78	82
VFI	78	76	64	70
M5Prime	83	85	76	80
Naive Bayes	82	81	74	72
AdaBoots M1/C4.5	80	81	80	78
AdaBoots M1/PART	80	83	79	78

[0119] In the above, the abbreviations are as follow: (LPF)sp.1 = low-passed signal of speaker 1 voice samples, (LPF)sp.2 = low-passed signal of speaker 2 voice samples, (SF)sp.1 = standard, non-filtered signal of speaker 1 voice samples, and (SF)sp.2 = standard, non-filtered signal of speaker 2 voice samples.

[0120] This table shows that use of only the quartiles of the low-passed signal intensity, still outperforms the combination of features used traditionally.

[0121] In order to attain the object of finding a very efficient small set of features, the Applicant implemented an automatic search method based on the use of genetic algorithms. Populations of features (limited to 30) were generated and evolved using as fitness the 10-fold cross-validation with two algorithms : Naive Bayes and 5-Nearest Neighbours (chosen mainly because they are fast to train). The exact genetic algorithm is the simple one described in Goldberg, D.E. (1989) "Genetic algorithms in search, optimization and machine learning, reading", MA: Addison-Wesley.

[0122] The outcome of this experiment was not obvious : within the selected feature set, there were features related to the quartiles of low-passed signal intensity and features related to the quartiles of pitch, but also features with relatively low individual information gain: those related to the quartiles of the minimas of the unfiltered smoothed intensity curve.

[0123] Also, it is again noted that the machine learning algorithm tends to always neglect features related to the variance or the range of distributions, whatever the measurement.

[0124] A final experiment using these 15 features along with all learning algorithms was conducted (max, min, median, 3rd quartile and 1st quartile of low-passed signal intensity, pitch and minimas of unfiltered signal intensity).

[0125] The results are summarised in table V below.

Table V

results of learning scheme using the "optimal" feature set		
name	speaker 1	speaker 2
1-NN	87	92
3-NN	90	92
10-NN	87	91
Decision Tree/C4.5	85	92
Decision Rules-PART	86	93
Kernel Density	87	91
Kstar	86	90
Linear Regression	83	89
LWR	87	89
Voted Perceptrons	65	78
SVM1	87	91
SVM2	90	96
SVM3	89	94
VFI	83	92
M5Prime	88	95
Naive Bayes	89	93
AidaBoots M1/C4.5	90	96
AidaBoots M1/PART	90	96

[0126] It can be observed the results obtained are very similar to the best results initially obtained, but with over than 10 times fewer features. Moreover, and interestingly, the variation between learning schemes is less important and algorithms which performed badly, like nearest neighbours or Naive Bayes, now behave more satisfactorily.

When only very few examples are provided

[0127] In the previous section, large training databases were used: this was optimal to explore feature and algorithmic spaces. However, when dealing with a speaker dependant task, this is not directly applicable to a real world robotic pet, for instance. Indeed, it is not conceivable that the owner of such a robot would give hundreds of supervised examples to teach it how to recognise its way of expressing basic emotions, even though this is what probably happens with human babies and real pets. Indeed, people are generally less inclined to spend such an amount of time with robotic pets.

[0128] It then becomes natural to ask what the results become if only very few training examples are given.

[0129] To this end, the Applicant produced an experiment using the "optimal" feature set fed earlier, but giving to each algorithms only 12 examples of each class. The algorithms were tested on the remaining items of the database.

[0130] This experiment was repeated 30 times with different sets of 12 examples, and results were averaged (the standard deviation was rather low, typically around 1.1).

[0131] Table VI summarises the results of the experiment.

Table VI

results of learning scheme when very few training examples are provided		
learning schemes	speaker 1	speaker 2
1-NN	81	85
3-NN	78	74

Table VI (continued)

results of learning scheme when very few training examples are provided		
learning schemes	speaker 1	speaker 2
10-NN	68	64
Decision Tree-C4.5	79	78
Decision Rules-PART	79	74
Kernel Density	82	84
Kstar	76	68
Linear Regression	75	74
LWR	75	79
Voted Perceptrons	40	51
SVM degree 1	81	83
SVM degree 2	83	85
SVM degree 3	83	85
VFI	77	76
M5Prime	80	79
Naive Bayes	85	87
AdaBoots M1/C4.5	79	84
AdaBoots M1/PART	79	90

[0132] It can be observed that some of the algorithms manage to keep a very reasonable level of performance (around 85 percent of success in generalisation). Among them are examples of very cheap algorithms like 1-nearest neighbour or Naive Bayes. These results are rather comparable (and in fact slightly superior) to what is described in Breazal, 2000 cited supra. However, in this case, learning was off-line with a larger database of several female speakers: what is important is that Breazal conducted experiments and showed that this level of success is sufficient to develop interesting interactions with a robotic pet. Also, Breazal showed how these results could be substantially improved when integrated into a larger cognitive architecture which is working in the real world. For example, linking this-recognition module to an-artificial limbic/emotional system in which there is some kind of emotional inertia (one very rarely swap from happy to angry in half a second) might give some additional information or tell the system there is uncertainty about the result. As a consequence, the robot may for instance take a posture showing it is not sure of what is happening. The human then will often repeat its utterance with an even more exaggerated intonation.

Conclusions on overall studies

[0133] From the foregoing, it appears that with the right set of features, reasonable performance can be reached when only few examples are given, which might be the case in "real situation" robots.

[0134] An aim is therefore to try to reduce the number of measurements while maintaining an acceptable level of performance, or possibly better. Indeed, some learning algorithms are not well suited to exploiting a large number of measurements, especially when are not all highly relevant. Some learning algorithms are of the constructive type. They are interesting inasmuch as they aim to generate a certain number of rules readable by the developer for classifying emotions. For example, with one or two databases, there is a learning algorithm which identified from the mass of measurements that is there is simply made a test on four chosen values, acceptable results can be obtained.

[0135] This allows to discern a certain number of things which appear to be important. These are typically the value of the signal linked to the intensity of the low-passed signal. More particularly, it indicates whether the median of the intensity of that signal is greater than a given value (e.g. 25.75 dB).

[0136] An experiment was repeated, by supplying to the learning algorithm only the different quartiles of the intensity of the low passed signal. Measurements show that the best learning algorithms achieve typically 80- 85% success (compared with 90-95% for a complete analysis). However, the calculation is greatly simplified, notably by the fact that no pitch analysis is conducted.

[0137] On the data base, experiments were made using values traditionally employed in the literature for comparison with either all the values or with the low passed signal intensity. It appears that using the same data base that the values obtained from known choices yield a success rate typically 70-75%. This is around 10% below the results obtained simply by using the low passed signal intensity.

[0138] A certain number of experiments were conducted to determine the best compromise between a small number of measurements and the results. This involved generating a large number of measurement groups, e.g. using a genetic algorithm, and analysing the space. The conclusion was that the best results are obtained from the combination of the quartiles of the low passed signal intensity + quartiles of the pitch + the quartiles of the minima of the normal (unfiltered) signal, giving 15 parameters. These success rates in the region of 90-97%, which is substantially the same as the success rate obtained by taking the 120 parameters together.

[0139] The next question is to determine whether a robot can generalise relatively well from a relatively small initial number of examples. To this end, an experiment was conducted in which the algorithms were supplied with just twelve examples per class - giving 48 examples in total - to determine to what extent the results would be different. It was observed that differences among the different learning algorithms became more pronounced. However, there are a certain number of standard algorithms which still achieve a success rate in the region of 85-90%.

[0140] The general results of the study are illustrated by the bar graphs of figures 9 to 13, which compare all the different cited learning algorithms in terms of correct classification of emotion. For each algorithm, two results are indicated, corresponding to each of two respective speakers of the database 8. More particularly :

- figure 9 shows performance with only low-passed signal intensity, for the correct classification in generalisation from an initial sample. The features are just the low-passed (250 Hz cut-off frequency) signal intensity;
- figure 10 shows performance with so-called "Kismet" features (Kismet being a system developed by MIT) namely: pitch + intensity + length segments;
- figure 11 shows performance for the best combination of features : low-passed signal intensity + pitch + minimas intensity;
- figure 12 show performance when very few samples are available, with the correct classification in generalisation, based on 12 training examples/class. The features used are: low-passed signal intensity + pitch + minimas intensity; and
- figure 13 is a chart summarising the results in terms of the best scores per experiment for each of the preceding charts.

Practical Implementation

[0141] Figure 14 illustrates schematically how the findings in accordance with the invention can be implemented in an emotion sensitive system 60. The latter can be part of a robotic pet, humanoid, teaching device, interactive computer or the like. Typically, the system is centred on a microcomputer comprising all the necessary processing, data storage, driving and interfacing functions necessary for the emotion recognition process. As these units are in themselves known and readily adaptable to application requirements by the skilled person, they have been omitted from the diagram for the sake of simplicity.

[0142] The system 60 presents an audio input section for receiving voice data from which an emotion is to be detected. The voice data can be obtained from two different sources: a built-in microphone 62 to pick up environmental sounds and thus detect a voice utterance from a nearby person on which to perform emotion detection, as in a real-life situation, and a database of recorded utterances, which can interface with the system 60 via an audio connector 64. The database can simply be the database 8 of figure 3, or an adapted version thereof. In this way, the system can be pre-trained, corrected, characterised, etc. This interfacing possibility is also useful to set up all the internal sound processing parameters, including the low-pass filtering characteristics.

[0143] The voice signals from either of the two sources 62 and 8 are fed to an audio amplification and preprocessing unit 66 which adapts the signal levels to the system electronics. The unit 66 may also include some noise reduction stages to reduce background noise to a minimum when the microphone 62 is in operation. The audio output 68 of unit 66 is fed along first and second paths P1 and P2 to two respective first and second channel inputs IC1 and IC2 of a feature extraction unit 70.

[0144] The first path P1 contains a first digital-to-analog converter (DAC) 72 between the audio output 68 and the first channel input IC1. The latter then simply receives the suitably digitised form of the amplified and preprocessed audio signal.

[0145] The second path P2 contains a low-pass filter 74 whose input is connected to the audio output 68. In the example, the low-pass filter is similar to the filter 18 of figure 3, inasmuch as it has an adjustable cut-off frequency. Typically, the cut-off frequency (Fco) is set by default to pass signals from a frequency range of 0 to 250 Hz, with a rapid attenuation after that frequency. The cut-off frequency can be set to different values by means of a cut-off fre-

quency selector 76. The latter can be manually controlled by means of a slider to produce different cut-off frequencies, and/or electronically adjusted by internal control circuitry, e.g. in feedback system to obtain optimal results during a characterisation phase. The typical range of possible values for the cut-off frequency F_{co} is from 150 Hz to 400 Hz. The adjustment in the cut-off frequency can be necessary to optimise the system 60 to different speakers, languages, environments, varieties of emotions to detect, etc.

[0146] The low-pass filter 74 further includes a variable attenuation control device 78 which allow the slope of the frequency response curve (cf. figure 4) to be modified, e.g. by changing the order of the filter, to allow a further degree of freedom in optimising the low-pass filtering characteristics.

[0147] However, the filter 74 need not be variable, nor have an adjustable slope, and simpler embodiments can be envisaged with a fixed low-pass filter set at a predetermined cut-off frequency, e.g. 250 Hz.

[0148] As for the system of figure 3, the low-pass filter 74 can be implemented with analog circuitry (passive or active), or as a digital filter. In the latter case, the DAC 80 is no longer necessary, but a similar DAC would be provided at its input, unless the digital low-pass filter comprises its own digital-to-analog converter stage. In a practical implementation, the low-pass filter 74 and its control units 76 and 78 would be typically fully digital such that the control of filter characteristics, including of selection of cut-off frequency are accessible digitally via an interface of a PC based control unit at the heart of the system 60.

[0149] The low-passed signal output from the low-pass filter 74 (analog in the figure) is supplied to a second analog-to-digital converter 80 which supplies the suitably digitised low-pass filtered signal to the second channel input IC2 of the feature extraction unit 70.

[0150] The latter forms the first link in a processing chain whose function is to detect an emotion conveyed by inputted voice signals.

[0151] The first part of the processing chain comprises the feature extraction unit 70, which is programmed to extract a selection of features from the digitised signal, of which some are to be extracted from the voice signal after low pass filtering (at channel input IC2) in accordance with the invention.

[0152] The feature extraction unit 70 operates in a manner analogous to the feature measurement station 30 of figure 3. In particular, it selectively extracts a set of features taken from the 120 features described supra, i.e. obtained from the possible combinations of:

- i) the intensity of the signal after low-pass filtering (derived from the signal at channel IC2),
- ii) the intensity of the signal without low-pass filtering (derived from the signal at channel IC1), and
- iii) the pitch of the signal without low-pass filtering (derived from the signal at channel IC1).

[0153] A 10 Hz smoothing of the above signals is performed.

[0154] From that stored 10 Hz smoothed signal are computed each of the three time-varying signals i) to iii) mentioned above.

[0155] Each thus-computed signal i) to iii) is digitally sampled at e.g. 0.01 second intervals. Each sampling gives an absolute instantaneous value of the intensity (signals i) and ii)) or of the pitch (signal iii)).

[0156] For each sampled signal, there is thus obtained a sequence of sampled values, typically a few hundred, given that a phrase is a few seconds long. The thus-obtained sequences of values are stored in time order, starting from the value of the earliest part of the signal, so that the sequence of values expresses the time variation of signal i), ii) or iii) considered.

[0157] For each sequence obtained from each of these three signals: intensity after low-pass filtering, intensity without low-pass filtering and pitch without low-pass filtering, the feature extraction unit 70 can calculate the four series of values described above in connection with the feature measurement station 30 of figure 3, i.e.

1) The series of local minimas, by determining the values in the sequence that correspond to a minima and constructing a series formed just of these minima values of the sequence. The series is stored in time order, starting from the earliest minima. A minima in the sequence of values is identified as a member of the sequence whose value is less than that of its neighbour immediately before and immediately after in the sequence.

2) The series of local maximas, by determining the values in the sequence that correspond to a maxima and constructing a series formed just of these maxima values of the sequence. The series is stored in time order, starting from the earliest maxima. A maxima in the sequence of values is identified as a member of the sequence whose value is greater than that of its neighbour immediately before and immediately after in the sequence.

3) The series of durations (time intervals) between local (adjacent) extremas - minimas and maximas - (which models rhythmic aspects of the signal) by identifying in the sequence the values that correspond to a minima or a maxima. This is achieved by comparing each value with its immediately adjacent neighbours as for series 1) and 2) above. Then, the positions in time of the identified extremas are analysed to determine the time intervals separating successive extrema. These time intervals are set out in a series, starting from the interval separating

EP 1 256 937 A2

the earliest and next earliest extremas, to yield the series of time intervals.

4) The series itself, by forming a series from all the successive sampled values of the sequence, starting from the beginning of the sampled voice signal.

5 **[0158]** On the resulting 3×4 (= 12) possible series, the feature extraction unit can perform the following 10 statistical operations to produce corresponding features:

- 1) the mean,
- 2) the maximum (i.e. the overall highest value),
- 10 3) the minimum (i.e. the overall lowest value),
- 4) the difference between the maximum and the minimum,
- 5) the variance,
- 6) the median,
- 7) the first quartile,
- 15 8) the third quartile,
- 9) the interquartile range, defined as the difference between the third and first quartiles,
- 10) the mean of the absolute value of the local derivative.

20 **[0159]** Naturally, some of the above statistical features involve a re-ordering of the values in the series from a time developing series to a series sorted by increasing value, notably to determine the first quartile, median and third quartile.

[0160] All the data are initially normalised before calculating these statistical characteristics.

[0161] The choice of which of these 120 features are to be extracted depends on the algorithms to be used, and is programmable. In the embodiment, the feature extraction unit 70 can be programmed to extract any number of these features according to requirements. At least one of the features is derived from an intensity of a low-pass filtered voice signal on which an emotion recognition is to be performed.

25 **[0162]** As an example, the feature extraction unit 70 can be programmed to extract at least one the following five statistical features among the ten listed above:

- the first quartile,
- 30 - the median,
- the third quartile,
- the maximum, and
- the minimum.

35 **[0163]** This extraction can be performed for one or a mix of all three signals i), ii, and iii) listed above, giving up to 15 features that can be extracted by unit 70 for each voice signal from which an emotion is to be detected, at least one of the features being extracted from a low-passed intensity signal.

[0164] Also, the extraction unit 70 can be additionally programmed to extract any one of features listed 1 to 10 above from the pitch signal extracted from the voice signal after low-pass filtering. In this case, for the thus-obtained low-pass filtered pitch, there can be obtained the same four series of values listed 1 to 4 above, each giving rise to the ten statistical features listed 1 to 10 above. Thus, the extraction unit can further be programmed to extract any number up to $4 \times 10 = 40$ statistical features from the pitch signal, derived after low pass filtering by filter 74.

40 **[0165]** The feature extraction unit 70 can also be programmed to extract a reduced number of statistical features which correspond e.g. to at least one of the first six features listed in table III above, i.e.:

- 1: the median value of sampled low-pass filtered signal intensity values,
- 2: the mean value of low-pass filtered signal intensity values,
- 3: third quartile value of series of sampled low-pass filtered signal intensity values,
- 4: the first quartile value of series of sampled low-pass filtered signal intensity values,
- 50 5: the maximum value of series of sampled low-pass filtered signal intensity values, and
- 6: the maximum value of series of sampled low-pass filtered signal intensity values.

[0166] Finally, each of the signals listed i) to iii) and - as the case arises - the low-pass filtered pitch signal can be individually programmed to be smoothed, e.g. by 10 Hz smoothing (or some other value), or not smoothed.

55 **[0167]** The extracted features are temporarily stored in a data logger 82 where they can be accessed by other functional units.

[0168] The emotion sensitive system 60 is configured to operate selectively in one of two possible modes, these being respectively off-line training emotion recognition and on-line training emotion recognition.

[0169] In the off-line training mode, the training to detect emotions is conducted outside the device, e.g. using the evaluation system as described with reference to figure 3. In this case, the programming unit PU of figure 3 delivers a complete emotion detection algorithm ready for use. The algorithm can be of the type set out in example 1, for instance, or some quite different algorithm. The emotion recognition can be frozen in its parameters and programming (closed algorithm), or it can be adaptable to some degree to update its parameters and programming on the basis of user feedback response to correct errors or to coach in emotion detection.

[0170] In the on line training mode, there is no initial set emotion recognition algorithm. Rather, there is provided a teaching algorithm which analyses the different features extracted from the feature extraction unit 70 and user feedback, or user indications of the emotions that are to be detected. From this input, the teaching algorithm develops an specific on-line emotion recognition algorithm, which can be continually modified, updated and improved with growing user input and feedback.

[0171] In figure 14, each of these two modes has its own path of functional units: OFF-L and ON-L respectively for the off-line and on-line modes.

[0172] For the off-line mode, there is provided an off-line emotion recognition algorithm unit 84 as mentioned above. The latter inputs the features extracted by unit 70, as present in the data logger 82 in accordance with its parameterisation and programming, and determines from these extracted features an emotion in accordance with its own pre-determined settings. In this mode, the voice signals will normally only come from the microphone 62.

[0173] For the on-line training mode, there is first constituted a database 86 containing extracted features from which different choices can be investigated by a learning algorithm unit 88. To this end, the database 86 is connected to the data logger 82 so as to accumulate all the extracted features and thus build a corresponding collection of extracted features. A feedback (loop FL1) from the learning algorithm unit 86 to the feature extraction unit 70 allows the latter to be instructed to change or select new features from the possibilities, in order to meet new needs deriving the most appropriate emotion detection algorithm. The teaching algorithm unit 88 can include a program allowing a user to input voice messages with different emotions, or be coupled to receive data from the a database of voices, such as database 8.

[0174] Once the teaching algorithm unit 86 has developed the required on-line emotion recognition algorithm, the latter is loaded into an on-line emotion recognition algorithm unit 90. The latter performs the same emotion recognition as that of unit 84, but is produced internally and specifically tailored to a given speaker, such as the user.

[0175] The types of emotion detected by either of the emotion detection algorithm units 84 or 90 can be as discussed earlier : calm, sad, happy, angry. Other emotions can of course be accommodated.

[0176] The choice of operation in the off-line or on-line is determined by a selector 92 which selects from which emotion detection algorithm 84 or 90 the detected emotions are to be taken. The selection is through an external user input signal 94.

[0177] Naturally, the following variants are possible:

- the system 60 is provided only with an off-line emotion recognition algorithm 84. The latter can then be totally rigid, or adaptable to a certain degree through feedback response, or
- the system is provided only with an on-line emotion algorithm system composed of the teaching algorithm and the on-line emotion recognition algorithm produced by the latter (units 86 to 90).

[0178] The programming unit PU of figure 3 serves to supply the required information and code for the different functional units. This is achieved through a programming interface PI. In particular, the programming unit PU and its interface PI serve to program the feature extraction unit 70, load the off-line emotion recognition algorithm (unit 84), and/or the teaching algorithm (unit 88), as well as features to database 86 in order to provide additional samples or complementary data.

[0179] The output from one or the other of the off-line and on-line emotion recognition algorithm units 84 and 90 is delivered to a detected emotion input unit 94 where the recognised emotion is acknowledged. From that unit 94, the detected emotion is delivered to an emotion response unit 96 for governing all the intended emotional responses. For instance, in the case of a robotic pet, the emotional response unit 96 will command the operation of different electro-mechanical actuators, lights and audio output devices to simulate an appropriate response.

[0180] One response includes giving feedback to the human user 98 or any other controlling person or device to indicate the detected information. This is achieved through a detected emotion indicator 100 which delivers a visual, aural, electrical or mechanical signal indicative of the detected emotion. This makes it possible to check if the identification is correct or not and to take corrective measures as appropriate.

[0181] In particular, it allows the user 98 to express a feedback response to a feedback response unit 102 that produces the suitable feedback messages to the various units concerned: off-line emotion recognition algorithm unit 84 (if the latter allows some flexibility), teaching algorithm unit 88, and on-line emotion recognition algorithm unit 90.

[0182] From the foregoing, it is clear that the invention can be implemented with many different combinations of

features extracted from the signal, at least one feature being extracted from the signal after low-pass filtering.
[0183] The choice of algorithm for the emotion detection on the basis of these features can be made on an empirical basis. The algorithms cited in this specification are just some of many algorithms suitable for that purpose. It can also be envisaged to implement the invention with other algorithms, both for teaching and for emotion detection.

Claims

1. A method of detecting an emotion conveyed in a voice signal, comprising the steps of:

- extracting (70) a set comprising at least one feature derived from said signal, and
- processing (84, 90) said set of extracted feature(s) to detect an emotion therefrom,

characterised in that it comprises a step of low-pass filtering (74) said voice signal prior to extracting therefrom at least one said feature of said set.

2. Method according to claim 1, wherein said low-pass filtering step comprises filtering (74) the intensity of said voice signal with a cut-off frequency (Fco) substantially in the range of 150 to 400 Hz.

3. Method according to claim 2, wherein said cut-off frequency (Fco) is substantially in the range of 250 to 300 Hz.

4. Method according to claim 3, wherein said cut-off frequency (Fco) is substantially equal to 250 Hz.

5. Method according to any one of claims 1 to 4, wherein the feature extracting step comprises the steps of selectively extracting one or more among 160 features obtained from the following combination of parameters I, II, and II:

- Parameters I:

- i) the intensity of the signal derived after said low-pass filtering step (74),
- ii) the intensity of the signal without low-pass filtering,
- iii) the pitch of the signal without low-pass filtering, and
- iv) the pitch of the signal derived after said low-pass filtering step (74);

- Parameters II, from a sequence of values of a parameter I obtained over a duration of said voice signal:

- 1) a series of local minimas, whose members are values in said sequence that correspond to a minima,
- 2) a series of local maximas, whose members are values in said sequence that correspond to a maxima,
- 3) a series of durations (time intervals), whose members are time values separating local (adjacent) extremas, and
- 4) a series, whose members are all the successive sampled values of the sequence; and

- Parameters III, calculated from a series of parameter II:

- 1) the mean,
 - 2) the maximum (i.e. the overall highest value),
 - 3) the minimum (i.e. the overall lowest value),
 - 4) the difference between the maximum and the minimum,
 - 5) the variance,
 - 6) the median,
 - 7) the first quartile,
 - 8) the third quartile,
 - 9) the interquartile range, defined as the difference between the third and first quartiles,
 - 10) the mean of the absolute value of the local derivative;
- at least one of said features originating from: i) the intensity of the signal after said low-pass filtering step (74) or from iv) the pitch of the signal after low-pass filtering step (74) of parameters I.

6. Method according to claim 5, wherein said parameters I comprise:

- i) the intensity of the signal derived after said low-pass filtering step (74),
 - ii) the intensity of the signal without low-pass filtering, and
 - iii) the pitch of the signal without low-pass filtering,
- whereby said combination yields a total of 120 different possible features;
- at least one of said features originating from: i) the intensity of the signal after said low-pass filtering step (74) of parameters I.

7. Method according to claim 5 or 6, wherein said parameters III comprise the following five parameters:

- 2) the maximum (i.e. the overall highest value),
- 3) the minimum (i.e. the overall lowest vale),
- 6) the median,
- 7) the first quartile,
- 8) the third quartile,

whereby said combination yields a total of 80 different possible features with claim 5, and 50 different possible combinations with claim 6.

8. Method according to claim 7, wherein said extracting step comprises extracting at least one of the following possible features:

- the maximum (i.e. the overall highest value),
 - the minimum (i.e. the overall lowest vale),
 - the median,
 - the first quartile,
 - the third quartile,
- calculated from:
- the series of values of the intensity signal derived after said low-pass filtering of the voice signal,
 - the series of minimum values of the intensity signal derived without low-pass filtering the voice signal,
 - the series of pitch values of the signal derived after said low-pass filtering of the voice signal, and
 - the series of pitch values of the signal derived without low-pass filtering of the voice signal.

9. Method according to any one of claims 1 to 8, wherein said extracting step comprises extracting features derived from both said voice signal after said low low-pass filtering step (74) and from said voice signal without low-pass filtering.

10. Method according to any one of claims 1 to 8, wherein said extracting step comprises extracting features only from said voice signal after said low-pass filtering step (74).

11. Method according to any one of claims 1 to 10, wherein said at least one member of said set of features comprises:

- 1: the median value of the signal intensity values derived after low-pass filtering the voice signal,
- 2: the mean value of the signal intensity values derived after low-pass filtering the voice signal,
- 3: third quartile value of the signal intensity values derived after low-pass filtering the voice signal,
- 4: the first quartile value of the signal intensity values derived after low-pass filtering the voice signal,
- 5: the maximum of the signal intensity values derived after low-pass filtering the voice signal, and
- 6: the maximum value of the series of the signal intensity values derived after low-pass filtering the voice signal.

12. Method according to any one of claims 1 to 11, further comprising a step of smoothing a signal extracted from said voice signal prior to said extracting step.

13. Method according to any one of claims 1 to 12, further comprising a step of generating an emotion detecting algorithm using a teaching algorithm, said teaching algorithm exploiting said at least one feature extracted from a low-pass filtered voice signal.

14. Method according to claim 13, further comprising a step of programming (PU) an emotion detection device (60) with an off-line emotion detection algorithm (84) and/or an on-line emotion detection algorithm (90), whereby said

algorithms exploit at least on feature extracted from a low-pass filtered voice signal.

15. A device for detecting an emotion conveyed in a voice signal, comprising the:

- extracting means (70) for extracting a set comprising at least one feature derived from said signal, and
- processing means (84, 90) for processing said set of extracted feature(s) to detect an emotion therefrom,

characterised in that it comprises a low-pass filtering means (74) for low-pass filtering said voice signal prior to extracting therefrom at least one said feature of said set.

16. Device according to claim 15, wherein said low-pass filtering means (74) has a cut-off frequency (F_{co}) substantially in the range of 150 to 400 Hz.

17. Device according to claim 16, wherein said cut-off frequency (F_{co}) is substantially in the range of 250 to 300 Hz.

18. Device according to claim 17, wherein said cut-off frequency (F_{co}) is substantially equal to 250 Hz.

19. Device according to any one of claims 15 to 18, wherein the feature extracting means (70) comprises means for selectively extracting one or more among 160 features obtained from the following combination of parameters I, II, and II:

- Parameters I:

- i) the intensity of the signal derived after said low-pass filtering step (74),
- ii) the intensity of the signal without low-pass filtering,
- iii) the pitch of the signal without low-pass filtering, and
- iv) the pitch of the signal derived after said low-pass filtering step (74);

- Parameters II, from a sequence of values of a parameter I obtained over a duration of said voice signal:

- 1) a series of local minimas, whose members are values in said sequence that correspond to a minima,
- 2) a series of local maximas, whose members are values in said sequence that correspond to a maxima,
- 3) a series of durations (time intervals), whose members are time values separating local (adjacent) extremas, and
- 4) a series, whose members are all the successive sampled values of the sequence; and

- Parameters III, calculated from a series of parameter II:

- 1) the mean,
- 2) the maximum (i.e. the overall highest value),
- 3) the minimum (i.e. the overall lowest value),
- 4) the difference between the maximum and the minimum,
- 5) the variance,
- 6) the median,
- 7) the first quartile,
- 8) the third quartile,
- 9) the interquartile range, defined as the difference between the third and first quartiles,
- 10) the mean of the absolute value of the local derivative;

at least one of said features originating from: i) the intensity of the signal after said low-pass filtering step (74) or from iv) the pitch of the signal after low-pass filtering step (74) of parameters I.

20. Device according to any one of claims 15 to 19, wherein said extracting means (70) is set to extract features derived from both said voice signal after said low low-pass filtering means (74) and from said voice signal without low-pass filtering.

21. Device according to any one of claims 15 to 19, wherein said extracting means (70) are set to extract features only from said voice signal after said low-pass filtering means (74).

EP 1 256 937 A2

22. Device according to any one of claims 15 to 21, wherein said at least one member of said set of features comprises:

- 1: the median value of the signal intensity values derived after low-pass filtering the voice signal,
- 2: the mean value of the signal intensity values derived after low-pass filtering the voice signal,
- 3: third quartile value of the signal intensity values derived after low-pass filtering the voice signal,
- 4: the first quartile value of the signal intensity values derived after low-pass filtering the voice signal,
- 5: the maximum of the signal intensity values derived after low-pass filtering the voice signal, and
- 6: the maximum value of the series of the signal intensity values derived after low-pass filtering the voice signal.

23. Device according to any one of claims 15 to 22, further comprising means for smoothing a signal extracted from said voice signal prior to said extracting means (70).

24. Device according to any one of claims 15 to 23, wherein said processing means comprises an off-line emotion recognition algorithm unit (84) containing a preprogrammed algorithm.

25. Device according to any one of claims 15 to 24, wherein said processing means comprises an on-line emotion recognition algorithm unit (90), said device further comprising a teaching algorithm unit (88) for producing an emotion recognition algorithm in response to data produced said extracting means (70).

26. A voice signal processing device specifically for the device according to any one of claims 15 to 25, **characterised in that** it comprises low-pass filtering means (74) for supplying a low-passed form of said voice signal to said extracting means (70).

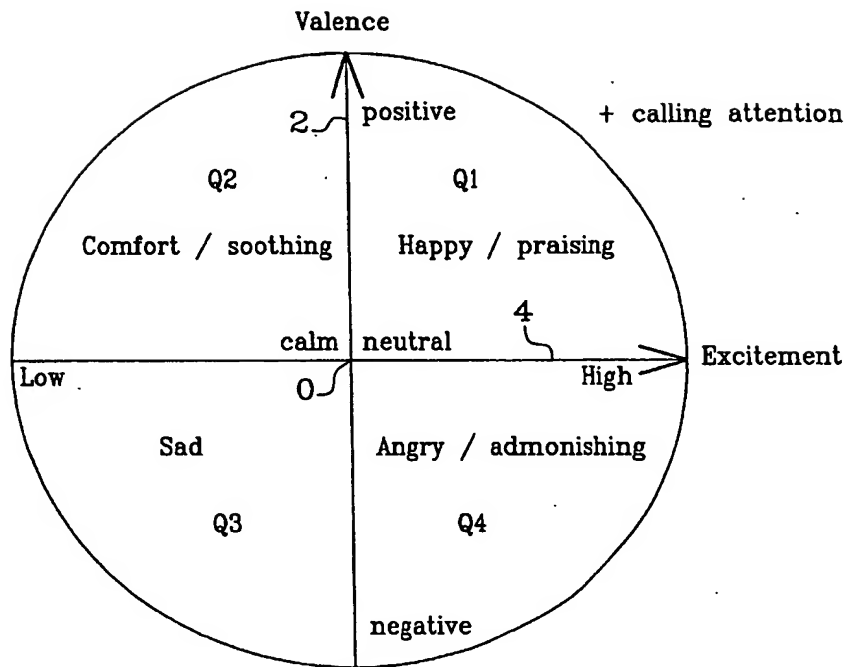


Fig. 1




Fig. 2a




Fig. 2b




Fig. 2c




Fig. 3

Transferred
audio
signal level




Fig. 4

Transferred
audio
signal level



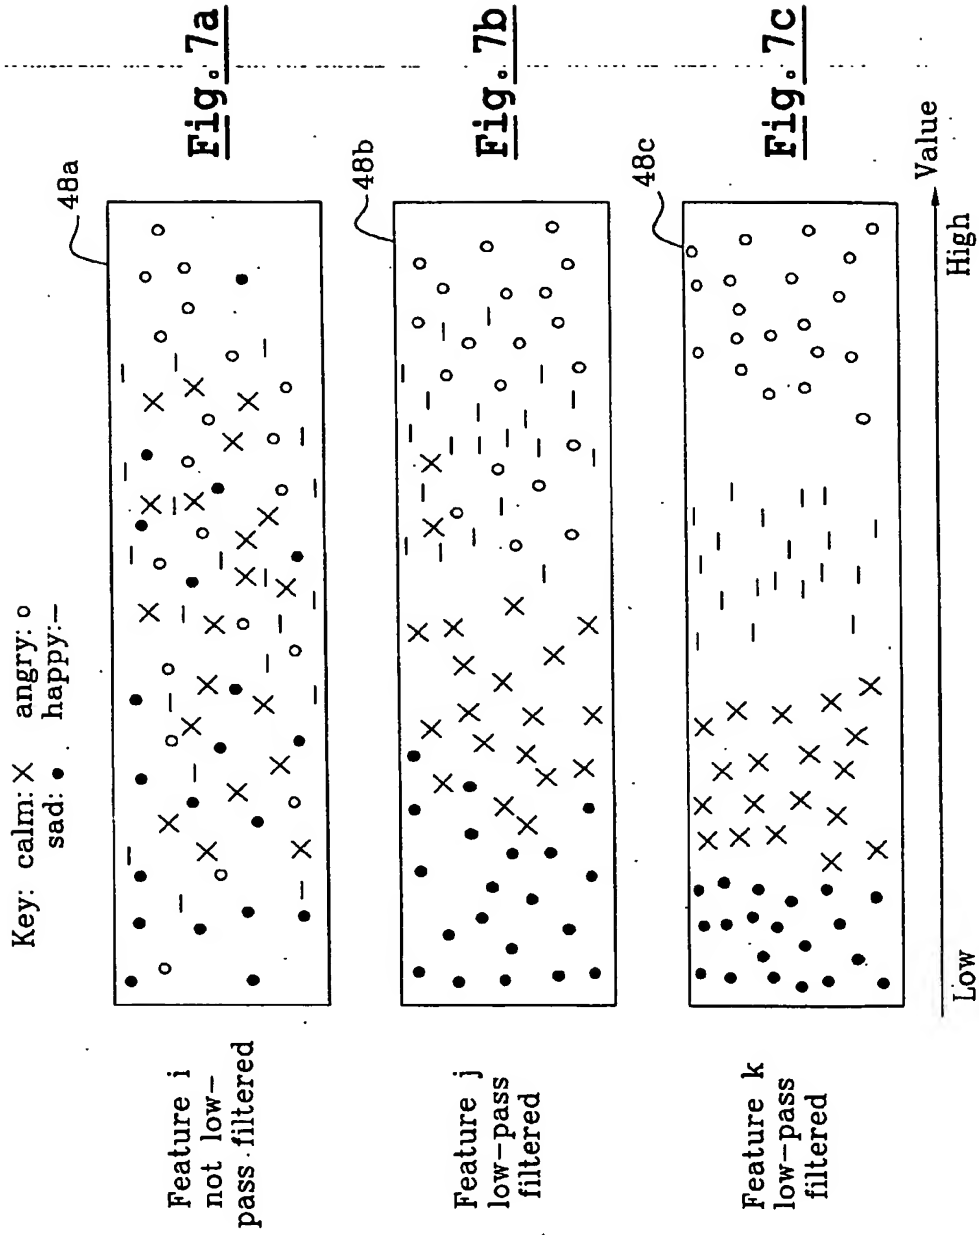



Fig. 5



1 measurement —→ 4 series of values

Fig. 6





PERFORMANCE WITH ONLY LOW-PASSED SIGNAL INTENSITY

Correct classification in generalisation, features=low-passed signal intensity




Fig. 9

PERFORMANCE WITH STANDARD FEATURES: PITCH+INTENSITY+LENGTH SEGMENTS

Correct classification in generalization, features=like standard: pitch+intensity+length segments

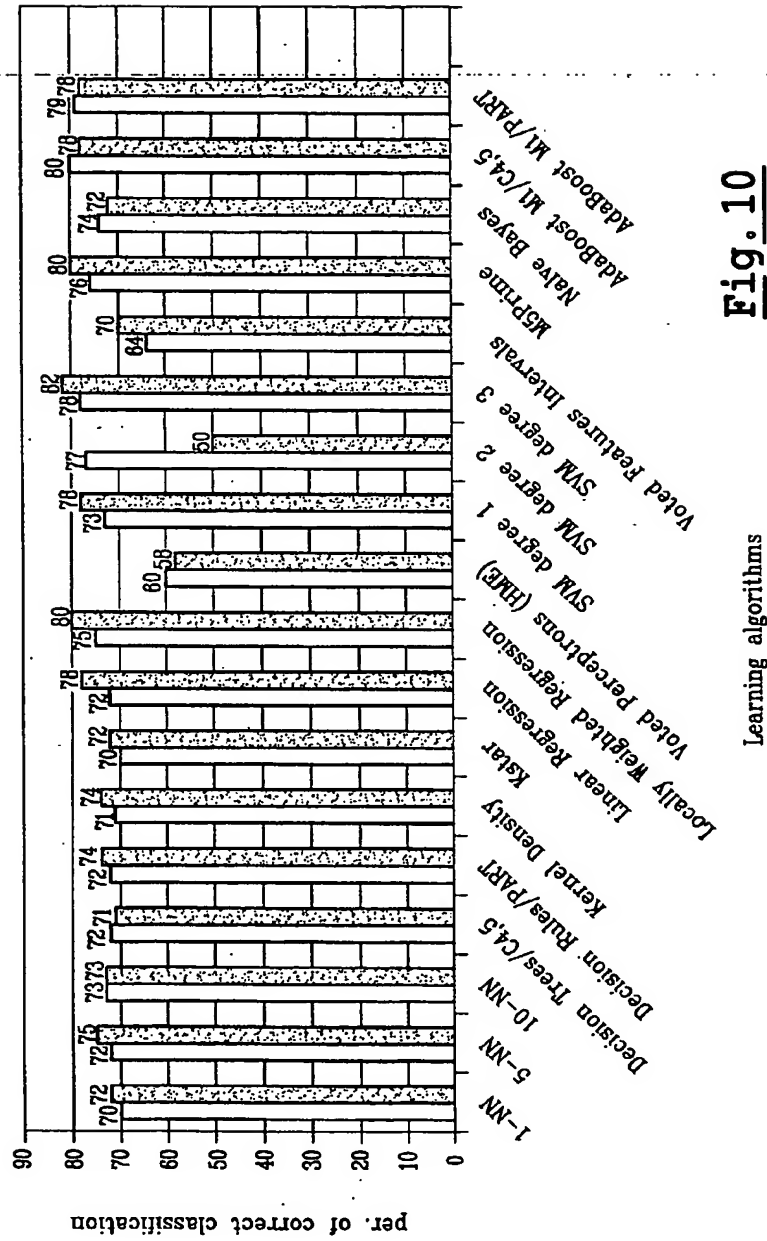


Fig. 10

Learning algorithms

PERFORMANCE WITH BEST COMBINATION OF FEATURES: LOW-PASSED
SIGNAL INTENSITY+PITCH+MINIMAS INTENSITY

Correct classification in generalisation, features=low-passed signal intensity+pitch+minimas intensity

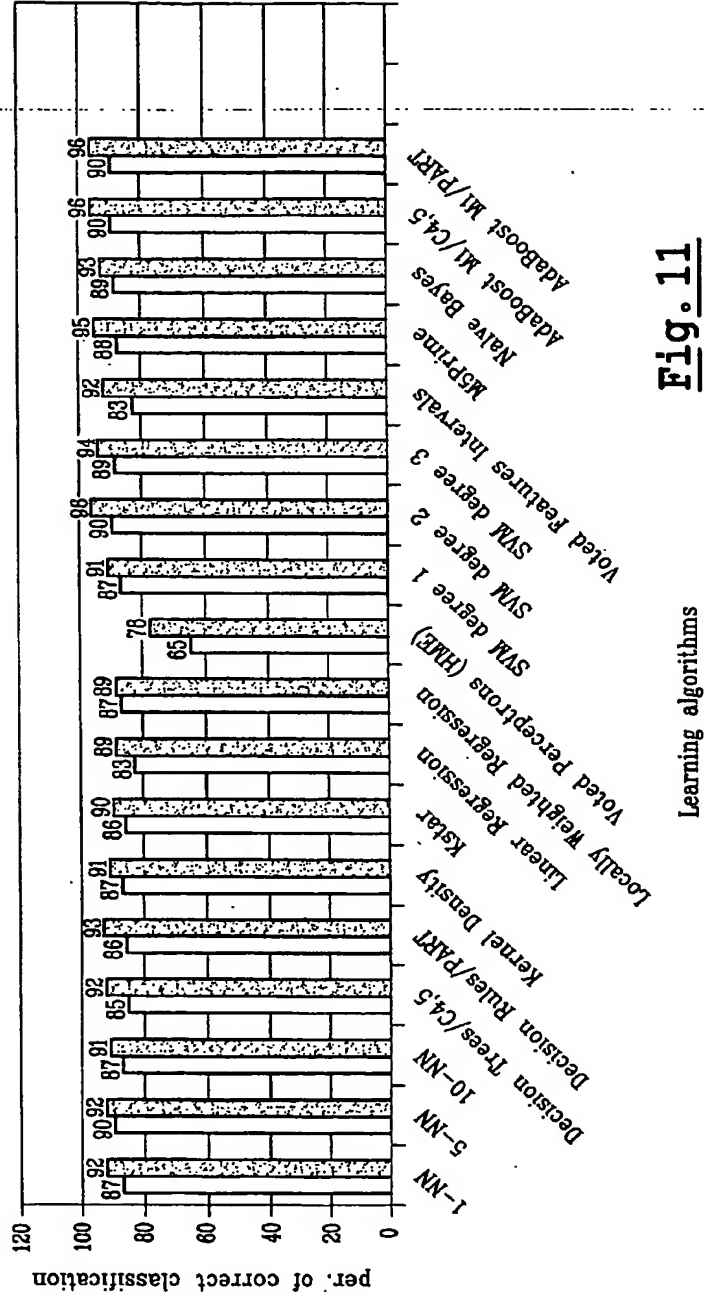


Fig. 11

Learning algorithms

WHEN VERY FEW EXAMPLES ARE AVAILABLE

Correct classification in generalisation, 12 training examples/class, features=
low-passes signal intensity+pitch+minimas intensity





Fig. 12

Learning algorithms

SUMMARY OF RESULTS
Best scores/experiment



Experiment

Fig. 13

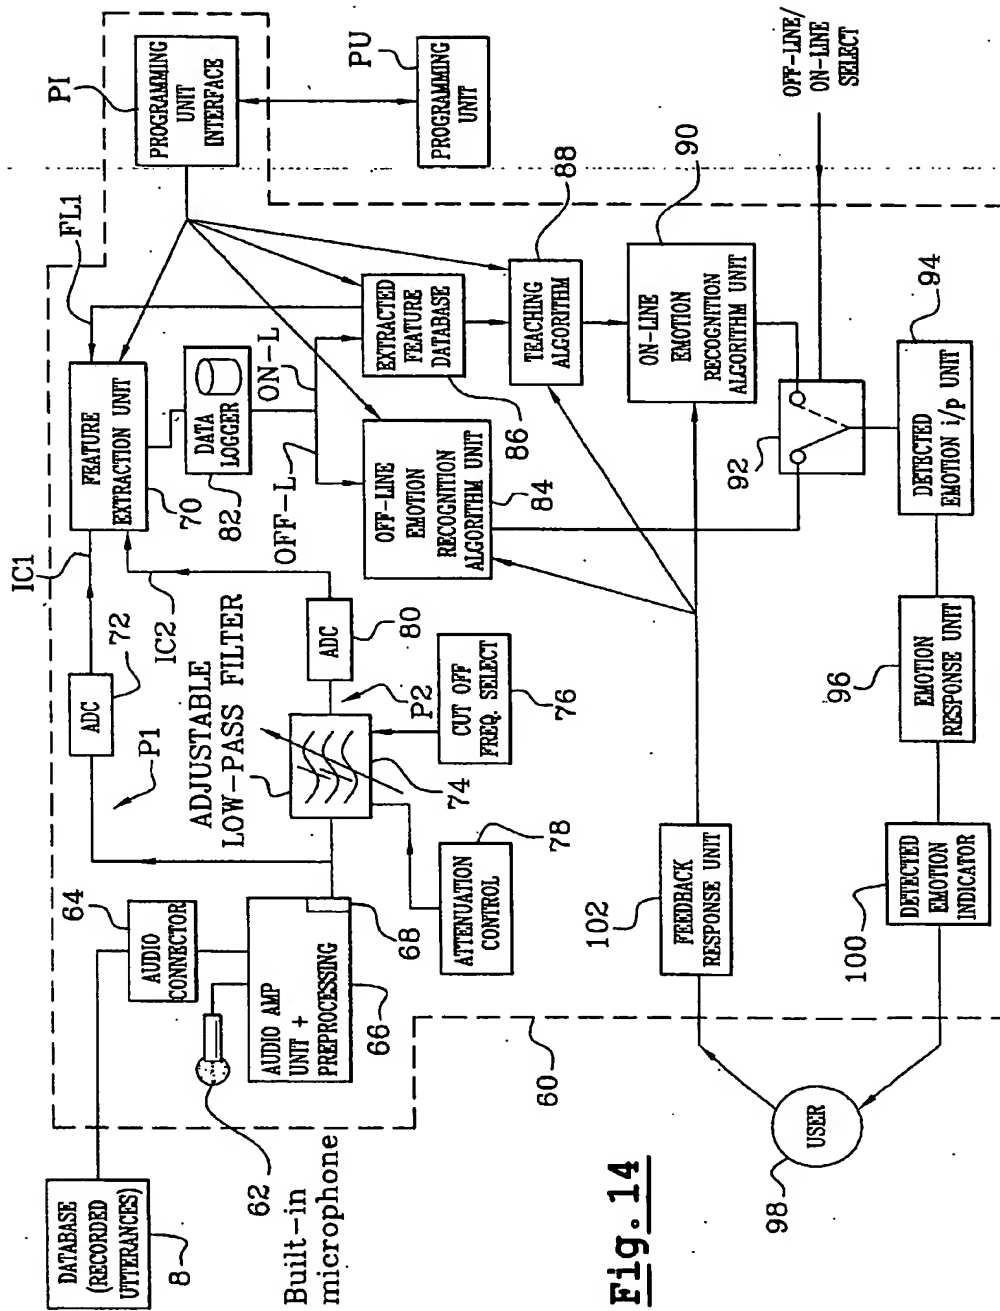


Fig. 14